High-Accuracy Asteroid Astrometry From Table Mountain Observatory

W. M. Owen, Jr. (wmo@jpl.nasa.gov), S. P. Synnott and G. W. Null
Jet Propulsion Laboratory, California Institute of Technology; 4800 Oak Grove Drive; Pasadena, California 91109-8099 USA

Abstract. We have installed a large-format CCD camera on the 0.6-meter telescope at JPL's Table Mountain Observatory and used it to obtain high-accuracy astrometric observations of asteroids and other solar system targets of interest. The detector contains 4096× 4096 15-μm pixels and produces a field of view 21′.9 square at a resolution of 0″.321/pixel. The field is wide enough to obtain sufficient reference stars from the ACT Catalog, and the oversampled images can produce centroids to a precision of 10–30 milliarcseconds. We describe the calibration and reduction procedures, including routine use of the overlapping field technique, that are necessary in order to obtain an accurate mapping of distortions in the focal plane. Our observations have enabled a successful flyby of 253 Mathilde, predicted ground tracks for several asteroid occultations, and will support onboard autonomous navigation for the Deep Space One mission.

Keywords: astrometry, asteroids, CCD, reduction techniques

1. Introduction

The current emphasis on robotic missions to small solar system bodies—minor planets such as 951 Gaspra, 243 Ida, 253 Mathilde, and 433 Eros; Europa and the other Galilean satellites of Jupiter; Pluto and Charon—has produced a need for high-accuracy astrometry of these bodies so that reliable ephemerides may be calculated. In order to meet this need, we have developed a wide-field CCD camera and associated observing techniques and reduction procedures that take advantage of the accurate Hipparcos and Tycho star catalogs (ESA 1997). The Hipparcos catalog, in particular, is extremely accurate but also relatively sparse, necessitating not only a wide field of view but also data reduction software that can produce an accurate map of the focal plane despite a paucity of reference stars. Eichhorn's (1960, 1988) overlapping plate method offers two distinct advantages in this regard. First, by synthesizing one effective field of view from a set of smaller exposures, one can work at lower reference star densities than otherwise possible. Second, the use of uncatalogued field stars enables us to obtain meaningful solutions for higher order terms in the reduction.

Our original intent was to use only Hipparcos stars as reference. The ACT Catalog (Urban & Corbin 1997) has provided us with a denser

set of supplemental, slightly less accurate reference stars; using both Hipparcos and ACT provides us more opportunities for observing any specified target and makes the data reduction more robust.

This paper describes all aspects of the system, hardware as well as software, and presents the results of the first year of operational asteroid astrometry.

2. Hardware

The choice of a large-format CCD was made only after a detailed study of various methods of obtaining Hipparcos-relative astrometry. A Ronchi ruling device (Gatewood 1987) was considered, but these are best suited for repeated observations of the same star field as opposed to solar system objects which move. CCDs operating in "scan mode," with the chip being read out while the images drift across the detector (Stone *et al.* 1996; Stone 1997), have shown great promise, but the limited time to build up the signal from each object provides not only a limiting magnitude for the system but also (Lindegren 1980; Han 1989) a lower bound on the astrometric accuracy due to short-term fluctuations in refraction. We were also concerned that low-frequency atmospheric effects might induce large-scale errors in the observations. A large-format CCD operated in sidereal or "stare" mode, however, can allow long exposures for faint objects or short ones for bright objects; temporal variations in refraction are more nearly correlated over the field.

Our CCD was fabricated by Loral North (now Lockheed Martin Fairchild Systems). The chip has 4096×4096 15-μ pixels, so that the active surface is 6.1 cm square. There are four read amplifiers, one in each corner, that operate in parallel in order to minimize the readout time. (Even so, to read out the entire 16 million pixels requires over 80 seconds.) The read noise is 12 electrons, and the full well is about 100 000 electrons per pixel, so that the dynamic range for a minimum S/N of 7 amounts to over 7.5 magnitudes. The gain is approximately 6 electrons per output data number (DN). The chip is front-illuminated, unthinned, and cooled with liquid nitrogen. We have procured three filters: Bessell (1990) R and I, and a narrowband (5 nm FWHM) interference filter centered at the 893 nm methane line.

This camera is used on a 24-inch (0.6-meter) $f/16$ Ritchey-Chrétien telescope at JPL's Table Mountain Observatory, located at an elevation of 2285 m in the eastern San Gabriel range in southern California. The telescope's actual focal length of 9452 mm yields a field of view $21'.9$ square at a resolution of $0''.321$/pixel. Star images are thus ad-

equately sampled for subpixel centroiding, while the field of view is large enough to capture, on average, 0.4 Hipparcos star and 3 ACT stars. Observations are therefore attempted only when the star background is significantly richer than average. At the very least, the target must lie inside a triangle formed by three accessible reference stars (or groups of stars). These stars need not lie within the same field of view, however, as the overlapping plate method is easily powerful enough to give good results even if the three stars appear once each on three different exposures.

3. Observing techniques

No matter how bright the target asteroid is, or how the reference stars are distributed around it, we always take at least two or three exposures of each target, using a different field for each exposure. Even if there are adequate reference stars in one field of view, additional offset exposures help to determine the higher-order field parameters as descibed below. In the more usual case, offset exposures are required anyway in order to capture enough reference stars.

We have found that the R filter yields better results than the I filter, even though the longer wavelength of the I passband promises better seeing. The I filter also has a wider passband, and the signal builds up more quickly, necessitating shorter exposures; the exposure time appears not to be long enough to allow the atmospheric seeing to be reduced to an acceptable level. Consequently most of our observations have been made through the R filter.

Most of the asteroids on our observing program are between ninth and fifteenth magnitude, and these form usable images without overexposing the reference stars, which themselves are usually eighth to tenth magnitude. A three-minute exposure through the R filter suffices for these targets (cf. Zacharias 1997) as it is both short enough to avoid long trailed images for mainbelt asteroids and long enough to average down the atmosphere.

A few of our target asteroids are considerably fainter, in some cases as faint as magnitude 19. These asteroids are observed in a two-step process. First, we observe the reference stars through the R filter, with at least three exposures as described above. Second, we remove the filter and take a series of long exposures, up to 20 minutes duration, with the target nearly centered in the field of view. For these long exposures we often command the telescope to remove half of the asteroid's apparent angular velocity, so that the asteroid and the stars trail at equal but

opposite rates. Again we offset the field between exposures, to help determine the plate parameters for the unfiltered frames.

It can also happen that some of the reference stars are overexposed, to the point of charge bleeding down columns, in a three-minute R exposure. When the observer detects an overexposed star image, he follows that exposure with a set of five to seven shorter exposures, of typically 15 to 20 seconds duration. These are placed onto the same CCD picture; the shutter opens and closes repeatedly before the chip is read out, and the observer slews the telescope while the shutter is closed. Thus one file contains multiple exposures, and the chip is read out but once. (The same technique is used to determine optimum focus, except that the telescope focus is changed between exposures as well.)

Likewise, a few of our targets (notably 1 Ceres and 4 Vesta) are considerably brighter than ninth magnitude. These asteroids would saturate the detector in just a few seconds, too short an interval to permit the atmospheric seeing to be averaged down to an acceptable level. For these targets we use the narrowband "methane filter" originally intended for observations of satellites of the outer planets. A three- to five-minute exposure through this filter suffices to bring up the reference stars and an adequate number of field stars without saturating the target.

At least once a night we take a sequence of calibration frames designed to determine the coefficients of the higher-order terms in the plate model; these are discussed in detail below.

We have found it unnecessary and a waste of both observing and processing time to perform the standard flat-field calibration procedures on our images. Slight pixel-to-pixel variations in sensitivity will increase the error in the computed locations of star images, but so will the additional noise provided by the flat-fielding process. Systematic variations in sensitivity, whether due to vignetting in the telescope optics, problems in the chip, or some other cause, will be manifest equally in the calibration frames and in the frames of target bodies. Therefore we expect that the calibration procedure will largely remove their astrometric effects.

4. Software

There are three main steps in the conversion of CCD frames into reduced positions of the target bodies which appear in them. First, the (x, y) centroids of the images must be found. Second, the images must be properly identified with the stars or asteroids whose light pro-

duced them. Last, the properly identified measured coordinates must be reduced to right ascension and declination.

4.1. Image centroiding

Much has been written over the years on various techniques for finding the centroid of a digital image (*e.g.*, Stone *et al.* 1996; Chui 1977). Our experience (Null *et al.* 1993) has been that fancy point-spread functions which do an excellent job of matching the observations do not necessarily produce precisely repeatable centroids. As long as the fitting function is circularly symmetric and the star images are reasonably symmetric as well, almost any simple fitting function will give good results. We have used several such fitting functions, and finally adopted a circular Gaussian profile above a constant background. This function has five parameters: the height and one-sigma width of the profile, the x and y coordinates of the center, and the background level. We obtain the centroid in two steps. First, we hold the width constant at a value representative of the seeing and solve for the other four parameters. Once convergence is reached, we then add the width to the parameter list and repeat. Convergence is fairly rapid for legitimate star images; indeed, a failure to obtain a solution is usually an indication that the filter was not operating on a star image at all. The uncertainty in the coordinates of a star image, after adjustment to unit weight, is typically between 0.02 and 0.10 pixel.

Since the overlapping plate method makes use of uncatalogued field stars as well as the reference stars, it is important to obtain centroids of every usable image in each frame. We search through the frame for each local maximum and attempt centroiding, saving the results when the centroiding process converges. Occasionally a bright, nearly saturated star image will display two or more local maxima due to shot noise in its core. Multiple valid solutions are then produced, and it is a simple matter to find these and consolidate them statistically into one entry.

Occasionally a target asteroid will be too faint to be swept up by the above process, or its image may be blended with that of a field star, or a reference Hipparcos star may have a close companion. In cases like these we can perform the centroiding operation offline using the same algorithm but in a different program, performing a joint solution for doubled images if necessary, and add the results to the file of usable observations.

4.2. Image identification

There may be hundreds of usable images in each frame, of which only a handful are reference stars or target bodies. Furthermore, the

telescope pointing and camera twist angle are not precisely known *a priori*. The eye can tell at a glance which are the reference stars, but to register manually each image is tedious and time-consuming work. We have therefore developed a procedure which nearly always produces the correct identification of the reference stars.

First we predict which reference stars and target bodies will appear in, or just outside, the nominal field of view. JPL's operational optical navigation software (Riedel *et al.* 1990; Owen & Vaughan 1991) produces a file containing topocentric, apparent predicted (x, y) image coordinates. This file, the file of observations, and a small input file containing (among other things) the atmospheric pressure, temperature, and humidity are passed to our "Automated Matching Program" (AMP; Owen 1996).

AMP processes all frames of one target, sequentially, in one execution. It reads the list of predicted image locations, determines the separation (in pixels) between each pair of images, and sorts the separations. Then it reads in the list of observed image locations, and likewise calculates pairwise separations for the N brightest images. The value of N is input by the user—we usually set it to 50—and the implicit assumption is that the reference stars will generally be brighter than the field stars.

For each pair of field stars, AMP checks to see if their separation agrees within a user-specified tolerance with the separations between any pair of predicted images. If there are any agreements, AMP tentatively identifies the two observed images with the two predicted images and performs a standard four-constant solution for center, twist, and scale. AMP uses this tentative solution to map each of the other predicted images into (x, y) coordinates, and it then counts how many predicted images now have a counterpart in the brighter observed images. The tentative identification that produces the most matches is adopted as the correct one.

Occasionally the automated registration process fails, usually because nearly all of the reference star images were overexposed and did not produce usable centroids. Such failures are easy to diagnose, and one can invoke a manual registration option in AMP to correct them. One merely provides in the user input to AMP the identification number of one or two predicted images and the identification of the corresponding observed image(s); AMP will use this user-supplied identification in lieu of the search process described above. If only one match was supplied, AMP will not change the input camera twist angle, and usually the twist is well enough known that subsequent processing is unaffected. Similarly, if *all* the reference star images proved unusable, one can edit their approximate coordinates into the file of observations,

go through AMP normally, and discard these images during reduction. Such a picture is still usable if other frames exist in which the reference stars in question were correctly exposed.

Identifying the reference stars and target bodies is but the first of AMP's functions. Once that identification is secure, by either automatic or manual registration, AMP then computes approximate values for the right ascension and declination of each unidentified image, again using the four-constant model. It maintains an internal list of all such positions, so that images of the same stars can be recognized as such across a set of frames.

AMP also computes differential atmospheric refraction (Stone 1984) relative to the center of the frame, and it applies this correction to each observation, first (using the *a priori* pointing) during the identification process, and again (using the final pointing) for each image. We currently have no provision for atmospheric dispersion, because *a*) we are observing in the red part of the spectrum, where dispersion is less, and rarely below 30 degrees altitude, and *b*) we usually observe through only one filter and thus do not obtain color information.

AMP produces three files to be used in subsequent processing. The first, and most important, is the list of observed image centroids, rewritten to include the identification of each image and the refraction correction. Second is an *ad hoc* catalog of astrometric positions for every star and for each observation of every target body. (The stellar positions are taken from the reference catalog when possible or from the four-constant model otherwise. Positions of target bodies are the originally predicted astrometric positions, *i.e.*, corrected for light time but not stellar aberration.) The third file contains the three Euler angles that specify for each exposure the UTC of the midpoint, the velocity $\mathbf{V}$ of the telescope relative to the solar system barycenter, the apparent J2000.0 right ascension and declination of the tangent point, and a twist angle about the optical axis.

4.3. Data reduction

The data reduction is carried out in two programs (Owen 1996). One of them computes the expected image locations, forms the residuals, and computes partial derivatives of the expected image locations with respect to the parameters of the model; the other carries out the least-squares adjustment. Separating these two functions allows the coefficients of the condition equations to be computed only once, so that one may vary the model or delete blunder points without incurring unnecessary reprocessing.

4.3.1. *Astrometric model*

The model for converting an astrometric position vector

$$\hat{\mathbf{x}}(\alpha, 90^\circ - \delta) \equiv (\cos\alpha\cos\delta, \sin\alpha\cos\delta, \sin\delta)^{\mathrm{T}}$$

into pixel coordinates (x, y) deviates slightly from standard procedures (*e.g.*, Eichhorn 1974) so we present it in full here.

Let the apparent J2000.0 right ascension and declination of the tangent point be α_{c} and δ_{c} respectively. Let ϕ represent the camera twist angle, which in our 3–2–3 representation is measured clockwise in the image plane from the $-\xi$-axis to the direction toward celestial north. (The nominal value for ϕ for our camera is $+90^\circ$, indicating that north is down in the image.) We apply stellar aberration according to the Newtonian formulation; the error in ignoring special relativity is on the order of 1 mas (Stumpff 1980), absorbed in the coordinates of the tangent point; and the error in differential aberration over the field is entirely negligible. Then we transform the apparent vector into a coordinate system whose third axis is aligned with the telescope optical axis and whose first and second axes are parallel to the column and row directions of the detector. Using the notation of Eichhorn (1974),

$$\rho = \mathbf{R}_3(\phi)\,\mathbf{R}_2(90^\circ - \delta_{\mathrm{c}})\,\mathbf{R}_1(\alpha_{\mathrm{c}})\,[\hat{\mathbf{x}}(\alpha, 90^\circ - \delta) + \mathbf{V}/c].$$

Next, we project the vector onto the focal plane whose distance f from the origin is the effective focal length (mm), and we transform the distances in the focal plane into pixel coordinates using the mapping matrix K (pixels/mm) and the pixel coordinates $\mathbf{p}_0 = (2048, 2048)^{\mathrm{T}}$ of the optical axis. The effects of differential refraction, precomputed by AMP, are added here as well. This step produces the gnomonic projection and then converts the ideal coordinates into pixels:

$$\begin{pmatrix} \xi \\ \eta \end{pmatrix} = f \begin{pmatrix} K_x & 0 \\ K_{yx} & K_y \end{pmatrix} \begin{pmatrix} \rho_x/\rho_z \\ \rho_y/\rho_z \end{pmatrix} + \mathbf{p}_0 + \text{refraction terms.}$$

Note that K_{xy} can be set to zero without loss of generality, since the camera twist angle ϕ already controls the position angle of the ξ-axis. Therefore K_{yx} can be interpreted as representing the departure from orthogonality of the ξ- and η-axes of the detector.

Finally, we allow for general deviations from the ideal projection, which we model as products of Legendre polynomials in each coordinate. The arguments of the Legendre polynomials are simply the original ξ and η mapped onto the interval $[-1, +1]$:

$$\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} \xi \\ \eta \end{pmatrix} + \sum_{i,j} \begin{pmatrix} A_{ij} \\ B_{ij} \end{pmatrix} P_i\left(\frac{\xi - 2048}{2048}\right) P_j\left(\frac{\eta - 2048}{2048}\right).$$

Separating the camera orientation parameters from the second-order terms in x^2, xy, and y^2 allows us to distinguish between effects that change from exposure to exposure and those that should be constant.

4.4. Reduction techniques

The set of adjustment parameters is not entirely identical to the set of parameters of the astrometric model described above. For instance, only two of the three parameters f, K_x, and K_y are independent, for one actually observes only two quantities, the scale in the ξ and η directions. Furthermore, we have striven to express parameters in forms that are easier to interpret. Hence the adjustment parameters are these:

- The scale s in the ξ direction, in units of $''$/pixel: $s \approx 205\,265/(fK_x)$;
- The pixel aspect ratio a, defined as the ratio of the scales in the η and ξ directions: $a = K_x/K_y$;
- The angle θ between the η-axis and the true normal to the ξ-axis: $\theta = K_{yx}/K_x$;
- Small rotations ϵ_x, ϵ_y, and ϵ_z of the camera about its three axes;
- Coefficients A_{ij} and B_{ij} of the Legendre polynomials, provided that $i + j > 1$;
- Changes to the coordinates of each "object," in the form $\Delta\alpha \cos\delta$ and $\Delta\delta$ (Eichhorn 1971, 1985).

The lowest-order Legendre coefficients are omitted from the parameter list because the other parameters render them redundant. The zeroth-order terms A_{00} and B_{00} shift the entire field uniformly in the x and y directions, respectively, and are thus equivalent to the leading terms produced by the small rotations ϵ_y and ϵ_x, respectively. The first-order terms A_{10} and B_{01} produce scale changes and are therefore replaced by s and a. Parameters A_{01} and B_{10} tilt the x- and y-axes and can be replaced by ϵ_z and θ.

The word "object" above means, in practice, a reference star, a field star, or a single observation of a solar system target body.

Magnitude, color, and coma terms are also available in the model, but we have not employed them. Magnitude-dependent errors in the observed position are not expected to be important for CCD detectors whose response to light is nearly linear. The Ritchey-Chrétien optical design eliminates coma, and indeed we see no evidence of it in defocused images.

We solve for the camera rotation angles ϵ_i and for the scale s separately for each exposure; the other parameters are assumed to be constant over a set of exposures. Solving directly for the objects' coordinates provides naturally all the power of the overlapping plate method.

The least-squares reduction itself is a linearized Kalman (1960) filter in which the covariance matrix **P** has been factored into the form $\mathbf{U}\,\mathbf{D}\,\mathbf{U}^{\mathrm{T}}$, with **D** diagonal and **U** upper triagonal with ones along the diagonal (Bierman 1977). This factorization preserves the stability features of square-root methods with the advantage that the algorithm never needs to compute a square root explicitly.

The number of exposures and thus the number of parameters other than object position parameters are known in advance. There may be hundreds of objects, however, and it is impractical to maintain two parameters for each object throughout the reduction. It is much more effective to sort the observations by object, store the two parameters $\Delta\alpha\cos\delta$ and $\Delta\delta$ in the first two rows and columns of the covariance matrix, and perform a "white noise" (uncorrelated) update to the covariance (Bierman 1977) when a new object is encountered in the sorted observation list. At each update, the top two rows of the covariance are saved along with the current values of the parameters, and the top two rows are reset to the correct *a priori* covariance for the new object—the catalog uncertainties for reference stars, or an unconstrained, uncorrelated diagonal matrix otherwise.

The *a priori* covariance matrix $\mathbf{P}_0$ is largely specified by user input. In the absence of an existing calibration (as for reducing calibration observations), $\mathbf{P}_0$ is diagonal with reasonably large variances for most parameters; the analyst may wish to provide smaller sigmas on certain parameters to constrain them to reasonable values. If a previous calibration is being used, its final covariance matrix becomes the *a priori* covariance for the next run. In either case the first two rows and columns, corresponding to the object position, are set to zero, with P_{11} and P_{22} set appropriately.

After the last observation has been processed, the computer holds in memory the final values for the corrections to the position of the final object and to the remaining estimated parameters. The saved corrections to the positions of the other objects, however, reflect only those observations that had been processed to that point. The second step, therefore, is the application of a Rauch–Tung–Striebel (1965) smoothing algorithm, as modified for the **U**-**D** factorization by Bierman (1983). The smoothing process goes through the observations in reverse order, changing the saved values of the position corrections for each object according to the information accumulated after the then-current

values had been saved. The smoothing process thus produces, object by object, the same values for the objects' positions that would have been obtained if all the objects' position parameters had been included at once in an adjustment of much larger dimension.

Whenever any object is observed but once, with an unconstrained *a priori* covariance on its position, the filter accounts for any residual in the observation by correcting the position of the object; there is no information in the observation to improve the estimate of the other parameters. It is only when an object appears on more than one exposure, or when a reference star is being observed, that the parameters of the model will be subject to adjustment. To treat each observation of a target body as a separate object, as we do, forces each observation to be independent of the others and produces a different position for each observation.

By contrast, when a field star is observed more than once, the final position is the weighted mean of the observations, given that the influence of that star on the plate model is properly taken into consideration. The final position for reference stars, however, is neither the catalog position nor the position implicit in the observations alone, but rather a weighted mean of them.

One consequence of this procedure is that spurious observations are not likely to corrupt the solution. If, for instance, a cosmic ray signal in one frame somehow survives the tests designed to discriminate against it, its provisional position is unlikely to be matched against that of any other object. Thus it will be treated as a field star observed only once, and it will not affect the solution. If there should be a field star in the vicinity, the matching process in AMP may well identify the cosmic ray with the star; the large residuals for this "star" will stand out, and one can delete the obvious blunder point easily and generate a new solution.

When an asteroid is so faint that it can be observed only without a filter, as described above, the reduction requires an additional step. The nightly calibration frames are used to reduce the shorter observations of the reference stars in the usual manner. The positions of the field stars, which are usually merely a byproduct of the reduction, now form a secondary reference catalog nearly on the system of the ACT. We then reduce the unfiltered observations to the secondary reference catalog, without using an *a priori* calibration model.

5. Calibration

We have alluded above to special calibration exposures, taken nightly, and reduced separately from the others. Following Monet (1992) and Null *et al.* (1993), we take at least five exposures of an open cluster (such as M38 or NGC 225) containing many ACT stars, with the cluster centered in the first exposure and the remaining exposures offset approximately a quarter of the field toward each corner in an $\times$ pattern. High-frequency deviations from the ideal gnomonic projection are manifest as systematic trends in the pairwise separations of star images from one exposure to the next. Since there are typically several dozen ACT stars and several hundred field stars in each exposure, it is easy to determine the parameters a, θ, and higher-order A_{ij} and B_{ij} to a few hundredths of a pixel. This technique cannot determine the linear terms a and θ in the absence of an accurate reference catalog; field stars can provide information only for higher-order terms. Test sequences revealed that terms beyond fifth order are small and poorly determined, and so we routinely solve for only those coefficients with $i + j \leq 5$.

By contrast, the three exposures of a target body may have relatively little overlap, and the star field may be sparse, so that it is much more difficult to obtain the Legendre coefficients directly from the target body exposures. Some of these terms reach 1 pixel in amplitude, as there is significant third-order distortion in a Ritchey-Chrétien system. However, on occasion the star field and overlap pattern even for target body astrometry can be sufficient to allow the full set of parameters to be determined.

We have therefore adopted the practice of including the calibration information in every reduction of target images. The nightly calibration frames are reduced first, with loose *a priori* constraints on the distortion parameters, and the solution and its covariance are saved. When the observations of the target bodies are processed, the values of the parameters thus determined (except for scale) are incorporated into the model for predicting the expected image locations, and the filter likewise solves for corrections to these *a priori* values. The filter also uses the full covariance matrix from the calibration as an *a priori* constraint. This technique is equivalent to including all of the calibration observations in each reduction of a target body, since the input covariance matrix amounts to no more than another set of condition equations thrown into the adjustment. Consequently, if the target body exposures have no power in themselves to improve the calibration, these additional observations will not affect those parameters greatly, and the results of the calibration will be carried through with little

change into the final solution. Conversely, the information contained in a particularly dense star field behind the target body will be combined statistically with the information from the calibration, perhaps even superseding the latter. In both cases the final solution uses all the available observations to their fullest.

6. Results

Implicit in the reduction procedure is the assumption that only the overall image scale changes significantly from one exposure to the next; the values of the other parameters, determined by calibration, are assumed to apply (although subject to possible improvement) throughout a night. Repeated test exposures of Praesepe, taken at the same pointing over a period of four hours, tend to confirm this assumption, for the postfit residuals over the entire set of images was 0.033 pixel in each coordinate, not significantly worse than over subsets of images, and indeed comparable to our best operational astrometry. This test also verifies the procedures for handling refraction. A second test used repeated sets of calibration frames, to examine the behavior of the distortion model both within a night and across nights. Finally, we again observed Praesepe, which contains six Hipparcos stars, and in the reduction treated one of the stars as as target body and recovered its catalog position to an accuracy of 30 mas.

The first operational use of the camera was to support the encounter of the Near Earth Asteroid Rendezvous spacecraft with 253 Mathilde in June 1997 (Yeomans *et al.* 1997, Scheeres *et al.* 1998). Mathilde was barely visible in the pre-dawn sky in the month before encounter, and we obtained 24 usable observations of it on seven mornings, with the last observation only three days before encounter. These observations seemed to indicate that Mathilde's ephemeris was in error by over twice its uncertainty in the direction of its motion.

During this time we also supplied astrometry to support the occultation of a star by 170 Maria. Our seven observations on three consecutive nights moved the occultation path by more than ten times its own width, and the occultation was observed within the revised path. This initial success provided convincing evidence that the astrometry was reliable, and it gave us confidence to suggest a higher-than-usual weight for the Mathilde observations. The observations reduced the uncertainty in Mathilde's position at encounter to less than 30 km, and the spacecraft used the updated ephemeris to obtain high-resolution imaging of Mathilde's surface, imaging that would have been lost otherwise.

Since then we have concentrated our efforts on a set of 30 brighter "beacon" asteroids to be used by the onboard navigation system of the New Millennium Deep Space 1 (DS1) mission (Bhaskaran *et al.* 1998), with the goal of improving their ephemeris uncertainties to a few tens of kilometers. (A few asteroids that were to occult stars have also been observed as time permitted.) We obtained some 488 observations, whose residuals appear in Fig. 1. The cluster of points to the northeast belongs to 4 Vesta, and the same trend in the residuals is seen in contemporary observations from other sites. The scatter in the residuals, excluding Vesta and a few obvious blunder points, is less than $0''.08$ in each coordinate.

The scatter in the asteroid residuals is somewhat higher than we had hoped, since the camera performed better in the tests described above. Much of the degradation in the residuals may arise from poor placement of reference stars; for the Praesepe test case, the target star was surrounded by five well-determined Hipparcos stars, while some asteroid observations were obtained relative to three or even only two ACT stars of lesser quality. Evidence to support this claim may be found in the residuals: for each asteroid, the scatter of the nightly residuals about their mean is on the order of $0''.04$, and the nightly means themselves show a somewhat larger scatter to produce the overall scatter shown in the figure.

Part of the degradation can perhaps be ascribed to the motion of the asteroids, for the residuals are somewhat worse in right ascension than in declination. There is also a substantial amount of periodic tracking error in the telescope's right ascension drive, and the tracking error can interact with the asteroid's motion to produce trailed images which are neither uniformly bright nor correctly centered on the asteroid's true position. We are currently investigating methods of compensating for the tracking error.

The residuals as a function of asteroid magnitude appear in Fig. 2. The relatively widely scattered points at magnitude 9.0 are from one night's observations of 1 Ceres that were very short exposures through the R filter instead of longer ones through the methane filter. There is a slight trend toward larger residuals for fainter asteroids, as one would expect, but there is no obvious degradation in quality at the point (about magnitude 16) where we begin to use a secondary reference catalog to reduce unfiltered observations.

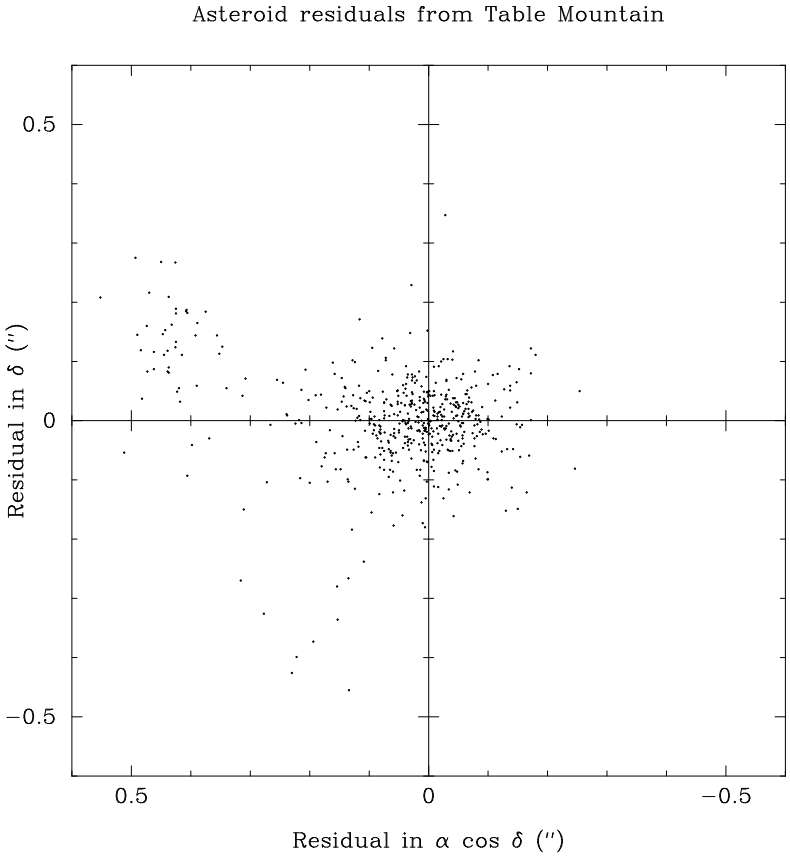

Figure 1. Residuals for observations of brighter asteroids.

7. Summary

We have described a large-format CCD camera whose wide field of view and small pixel size are yielding astrometry of asteroids to an accuracy better than 0″.1. To obtain results this accurate requires the use of the overlapping plate method during the reduction process, and

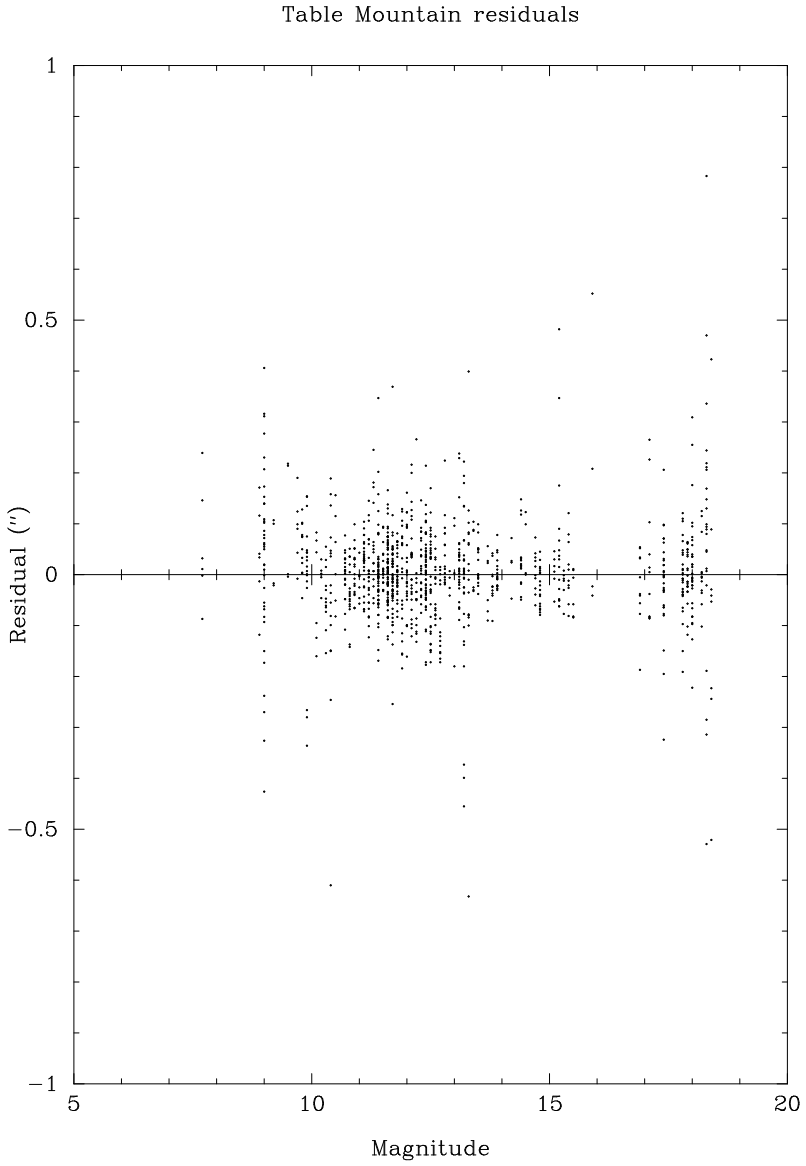

Figure 2. Distribution of residuals as a function of asteroid magnitude.

the observations themselves are made in such a manner as to bring the full power of that method to bear on the problem. Our work has contributed directly to the successful flyby of 253 Mathilde by the NEAR spacecraft, to the prediction of several asteroid occultations, and to the preparations for the DS1 mission.

Acknowledgements

The camera hardware and electronics were designed, built, and tested by Andy Collins, Tom Elliott, Walt Proniewicz, Luis Ramirez, David Erickson (all JPL), and Arsham Dingizian (Caltech). The camera control software was developed by Dale Winther (JPL). Steve Gillam and Jim Young, both of JPL's Table Mountain Observatory, provided invaluable onsite support and technical advice.

One of us (WMO) wishes particularly to acknowledge the influence of Heinrich Eichhorn, not only regarding this project (as evident from the references below), but in many other areas as well.

The work described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

References

Bessell, M. S. 1990, PASP, 102, 1181

Bhaskaran, S., Desai, S. D., Dumont, P. J., Kennedy, B. M., Null, G. W., Owen, Jr., W. M., Riedel, J. E., Synnott, S. P., & Werner, R. A. 1998, "Orbit Determination Performance Evaluation of the Deep Space 1 Autonomous Navigation System," Amer Astronautical Soc paper AAS 98–193, AAS/AIAA Space Flight Mechanics Mtg, Monterey, California

Bierman, G. J. 1977, Factorization Methods for Discrete Sequential Estimation (New York, Academic Press), chapter V

Bierman, G. J. 1983, Automatica, 19, 305

Eichhorn, H. K. 1960, AN, 285, 233

Eichhorn, H. K. 1971, Ann Univ-Sternwarte Wien, 30, 5

Eichhorn, H. K. 1974, Astronomy of Star Positions (New York, Frederick Ungar)

Eichhorn, H. K. 1985, A&A, 150, 251

Eichhorn, H. K. 1988, Proc IAU Symp 133 (Mapping the Sky), S. Débarbat *et al.*, eds. (Dordrecht, Kluwer), 177

European Space Agency 1997, The Hipparcos and Tycho Catalogues, ESA SP–1200 (Noordwijk, ESA Publ Div)

Gatewood, G. D. 1987, AJ, 94, 213

Han, I. 1989, AJ, 97, 607

Kalman, R. E. 1960, J Basic Engineering, March issue, 35

Lindegren, L. 1980, A&A, 89, 41

Monet, D. G. 1992, Wide Field/Planetary Camera Final Orbital/Science Verification Report, S. M. Faber, ed. (Baltimore, Space Telescope Science Inst), chapter 7

Null, G. W., Owen, Jr., W. M., & Synnott, S. P. 1993, AJ, 105, 2319

Owen, Jr., W. M. 1996, "Automated Astrometric Data Analysis System User's Guide," JPL Interoffice Memorandum 312.8–96–013 (internal document)

Owen, Jr., W. M., and Vaughan, R. M. 1991, "Optical Navigation Program Mathematical Models," JPL Engineering Memorandum 314–513 (internal document)

Perryman, M. A. C. 1997, private communication

Rauch, H. E., Tung, F., & Striebel, C. T. 1965, AIAA J, 3, 8, 1445

Riedel J. E., Owen, Jr., W. M., Stuve, J. A., Synnott, S. P., & Vaughan, R. M. 1990, "Optical Navigation During the Voyager Neptune Encounter," Amer Inst of Aeronautics & Astronautics paper 90–2877, AIAA/AAS Astrodynamics Conference, Portland, Oregon

Scheeres, D. J., Dunham, D. W., Farquhar, R. W., McAdams, J. V., Helfrich, C. E., Owen, Jr., W. M., Synnott, S. P., Williams, B. G., Wolff, P. J., & Yeomans, D. K. 1998, "Mission Design and Navigation of NEAR's Encounter with Asteroid 253 Mathilde," Amer Astronautical Soc paper AAS 98–184, AAS/AIAA Space Flight Mechanics Mtg, Monterey, California

Stone, R. C. 1984, A&A, 138, 275

Stone, R. C. 1997, AJ, 113, 2317

Stone, R. C., Monet, D. G., Monet, A. K. B., Walker, R. L., Ables, H. D., Bird, A. R., & Harris, F. H. 1996, AJ, 111, 1721

Stumpff, P. 1980, A&A, 84, 257

Urban, S. E. & Corbin, T. E. 1997, The ACT Reference Catalog (Washington, U. S. Naval Observatory)

Yeomans, D. K., Barriot, J.-P., Dunham, D. W., Farquhar, R. W., Giorgini, J. D., Helfrich, C. E., Konopliv, A. S., McAdams, J. V., Miller, J. K., Owen, Jr., W. M., Scheeres, D. J., Synnott, S. P, & Williams, B. G. 1997, Science 278, 2106

Zacharias, N. 1997, AJ, 113, 1925